MEHMET KEREM TURKCAN

Associate Research Scientist, Civil Engineering & Engineering Mechanics, Columbia University ◇ mkt2126@columbia.edu
mehmetkeremturkcan.com ◇ github.com/mkturkcan ◇ huggingface.co/mehmetkeremturkcan
orcid.org/0000-0001-9273-7293 ◇ scholar.google.com/citations?user=306TgWoAAAAJ

RESEARCH INTERESTS

My research is on physical AI: models that perceive, simulate, and act in the physical world, built for cities and medicine. It connects three lines of work.

- ◇ **Simulation and sim-to-real transfer.** Training perception and robots for real streets in simulation, and understanding what transfers. With the Constellation benchmark, published in IJCV, we showed that detectors trained on CARLA renders of the same camera fall far short on real images, and Boundless narrows that gap with photorealistic Unreal Engine 5 data. Valdrada extends this to a real-time digital twin of New York City compiled from public records, and G1 Patrol trains a Unitree G1 humanoid in Isaac Lab to navigate crowded sidewalks.
- ◇ **Foundation models for perception and action.** Vision-language models (VLMs) and open-vocabulary and monocular 3D perception that run in real time on edge devices, and vision-language-action (VLA) and world models for robots, including Open-H-Embodiment, the dataset behind the GR00T-H VLA for medical robotics, and world models of robotic suturing.
- ◇ **Language models and agents.** Robust post-training of large language models (LLMs), reinforcement learning of communication among small language models, generative agents, and language models on edge devices.

Earlier, I built connectome-scale models of the fruit fly brain. I release this work as open models, datasets, and research software.

ACADEMIC APPOINTMENTS

◇ **Columbia University**, New York, NY — Jul 2024 – Present
Associate Research Scientist, Civil Engineering & Engineering Mechanics
- Research in the NSF Engineering Research Center for Smart Streetscapes (CS3), a five-institution center; lead applied AI projects in physical AI, real-time urban perception, edge AI, vision-language models, urban digital twins, and medical robotics, including surgical world models.
- Lead work toward a civic robot system for CS3 applications: humanoid and quadruped navigation in dense pedestrian environments, trained in simulation with Isaac Lab.
- Main engineering instructor for the CS3 Research Experience for Teachers (2024–2026) and co-chair of the CS3 Student Leadership Council (2024–2025).

◇ **Columbia University**, New York, NY — Jan 2023 – Dec 2023
Lecturer in Electrical Engineering
- Designed and taught EECS E6691 Advanced Deep Learning and ECBM E4040 Neural Networks and Deep Learning to 123 graduate students; recognized with the 2024 Distinguished Teaching Award.

◇ **Columbia University**, New York, NY — Nov 2022 – Jul 2024
Postdoctoral Research Scientist, Electrical Engineering
- Built real-time digital twins of New York City intersections, synthetic data and benchmarks for dense urban perception, and computer vision for robotic surgery and endoscopy training with Northwell Health collaborators; mentored 11 Master's and three high school researchers.

EDUCATION

◇ **Columbia University**, New York, NY — 2022
- Ph.D. in Electrical Engineering; Systems Biology and Neuroengineering, GPA 4.10/4.33.
- Dissertation: *Sensory Processing and Associative Learning in Connectome-Based Neural Circuits.* Advisor: Aurel A. Lazar.

◇ **Columbia University**, New York, NY — 2016
- M.S. in Computer Science; Machine Learning and Thesis Track, GPA 3.97.

◇ **Istanbul Technical University**, Istanbul, Turkey — 2015
- B.S. in Electronics and Communication Engineering, GPA 3.75/4.00.

HONORS AND AWARDS

◇ **Best Paper Award**, ACM/IEEE Symposium on Edge Computing (SEC) — 2025
◇ Finalist, INRIX x MetroLab Challenge, for validating AI traffic models for New York City — 2025
◇ Distinguished Teaching Award, Columbia University Electrical Engineering — 2024

◇ Smart Cities North America Award, Transportation, IDC Government Insights, for the Columbia Digital Twin of NYC project 2023

◇ National Postdoc Appreciation Week Excellence Award, Columbia University 2023

◇ Herbert French Fellowship, Columbia University; Helmsley Fellowship, Cold Spring Harbor Laboratory 2017–2019

PUBLICATIONS

275 citations and an h-index of 12 on Google Scholar as of October 2026. Labels: [J] journal articles, [C] peer-reviewed conference papers, [W] workshop papers, demos, and short papers, [P] preprints, [A] conference abstracts, [D] dissertation.

Journal Articles

[J1] **Newell-Car-Following-Model-Informed Gaussian Process Regression Model for Ultra-Short-Term Arterial Vehicle Trajectory Prediction.** N. S. Ahmad, P. J. Jin, C. Zang, **M. K. Turkcan**, and Z. Kostic. *Transportation Research Record, to appear, 2026.*

[J2] **Harnessing Floating Car Data, Traffic Camera Observations, and Network Flow Analysis for Traffic Volume Estimation.** A. Kosikova, **M. K. Turkcan**, A. Darrat, and A. Smyth. *Discover Civil Engineering, 3(1):257, 2026.* doi.org/10.1007/s44290-026-00657-x

[J3] **Constellation Dataset: Benchmarking High-Altitude Object Detection for an Urban Intersection.** **M. K. Turkcan**, C. Zang, S. Narasimhan, G. H. Je, B. Yu, M. Ghasemi, J. Ghaderi, G. Zussman, and Z. Kostic. *International Journal of Computer Vision, 134(10):429, 2026.* doi.org/10.1007/s11263-026-03029-1

[J4] **AI-Powered CPS-Enabled Vulnerable-User-Aware Urban Transportation Digital Twin: Methods and Applications.** Y. Fu, **M. K. Turkcan**, M. Ghasemi, Z. Mo, C. Zang, A. Adhikari, Z. Kostic, G. Zussman, and X. Di. *IEEE Transactions on Intelligent Transportation Systems, 27(5):4949–4966, 2026.* doi.org/10.1109/TITS.2026.3656165

[J5] **Computer Vision Scoring of Endoscopically Traced Line Figures in an Inanimate Scope Tip Coordination Training Model.** N. Mitra, S. Narasimhan, **M. K. Turkcan**, Z. Kostic, Y.-R. Chen, S. Choksi, E. Nilsson-Sjolander, P. Palacios, K. Carsky, A. Riaz, F. Filicori, and R. L. Whelan. *Surgical Endoscopy, 40(1):657–664, 2026.* doi.org/10.1007/s00464-025-12354-4

[J6] **Monocular 3D Tooltip Tracking in Robotic Surgery: Building a Multi-Stage Pipeline.** S. Narasimhan, **M. K. Turkcan**, M. Ballo, S. Choksi, F. Filicori, and Z. Kostic. *Electronics, 14(10):2075, 2025.* doi.org/10.3390/electronics14102075

[J7] **Automatic Assessment of Robotic Suturing Utilizing Computer Vision in a Dry-Lab Simulation.** S. Choksi, S. Narasimhan, M. Ballo, **M. K. Turkcan**, Y. Hu, C. Zang, A. Farrell, B. King, J. Nussbaum, A. Reisner, Z. Kostic, G. Taffurelli, and F. Filicori. *Artificial Intelligence Surgery, 5(2):160–169, 2025.* doi.org/10.20517/ais.2024.84

[J8] **Severity of Coronary Artery Disease Is Associated with Diminished circANRIL Expression: A Possible Blood Based Transcriptional Biomarker in East Africa.** G. Akan, E. Nyawawa, B. Nyangasa, **M. K. Turkcan**, E. Mbugi, M. Janabi, and F. Atalar. *Journal of Cellular and Molecular Medicine, 28(3):e18093, 2024.* doi.org/10.1111/jcmm.18093

[J9] **Surgical Phase Recognition in Inguinal Hernia Repair: AI-Based Confirmatory Baseline and Exploration of Competitive Models.** C. Zang, **M. K. Turkcan**, S. Narasimhan, Y. Cao, K. Yarali, Z. Xiang, S. Szot, F. Ahmad, S. Choksi, D. P. Bitner, F. Filicori, and Z. Kostic. *Bioengineering, 10(6):654, 2023.* doi.org/10.3390/bioengineering10060654

[J10] **A Programmable Ontology Encompassing the Functional Logic of the *Drosophila* Brain.** A. A. Lazar, **M. K. Turkcan**, and Y. Zhou. *Frontiers in Neuroinformatics, 16:853098, 2022.* doi.org/10.3389/fninf.2022.853098

[J11] **Accelerating with FlyBrainLab the Discovery of the Functional Logic of the *Drosophila* Brain in the Connectomic and Synaptomic Era.** A. A. Lazar, T. Liu, **M. K. Turkcan**, and Y. Zhou. *eLife, 10:e62362, 2021.* doi.org/10.7554/eLife.62362

[J12] **Intronic Variants in the Long Non-Coding RNA CDKN2B-AS1 Are Strongly Associated with the Risk of Coronary Artery Disease in the Northern Tribes of Tanzania.** G. Akan, P. Kisenge, T. S. Sanga, E. Mbugi, **M. K. Turkcan**, M. Janabi, and F. Atalar. *Tanzania Journal of Health Research, 21(1):1–14, 2019.* doi.org/10.4314/thrb.v21i1.3

[J13] **Common SNP-Based Haplotype Analysis of the 9p21.3 Gene Locus as Predictor Coronary Artery Disease in Tanzanian Population.** G. Akan, P. Kisenge, T. S. Sanga, E. Mbugi, I. Adolf, **M. K. Turkcan**, M. Janabi, and F. Atalar. *Cellular and Molecular Biology, 65(6):33–43, 2019.* doi.org/10.14715/cmb/2019.65.6.7

[J14] **Generation of $1/f^{\alpha}$ Noise via Frequency Scaling. M. K. Turkcan** and T. Akgül. *EMO Bilimsel Dergi, 4(8):41–46, 2015.* dergipark.org.tr/en/download/article-file/63864

Peer-Reviewed Conference Papers

[C1] **Zero-Shot Image Tagging at Traffic Intersections via SAM-Guided Vision-Language Models.** W. Ho, **M. K. Turkcan**, J. Ghaderi, G. Zussman, and Z. Kostic. *International Conference on Machine Learning and Applications*

(ICMLA), pp. 243–250, 2026. doi.org/10.1109/ICMLA69890.2026.00042

[C2] **Digital Eyes on the Road: Using Street Cameras to Verify Traffic Integrity and Detect Sybil Attacks at City Scale.** J. Tavori, **M. K. Turkcan**, Z. Kostic, S. Stolfo, and G. Zussman. *USENIX Symposium on Vehicle Security and Privacy (VehicleSec), 2026.* usenix.org/conference/vehiclesec26/presentation/tavori

[C3] **Worst-Case Attacks in Reactive Edge-Cloud Systems: Expected Latency and SLA Violation.** J. Tavori, **M. K. Turkcan**, Z. Kostic, J. Ghaderi, and G. Zussman. *IEEE Conference on Computer Communications (INFOCOM), pp. 1–10, 2026.* doi.org/10.1109/INFOCOM59046.2026.11571607

[C4] **Real-Time Video Analytics for Urban Safety: Deployment over Edge and End Devices.** M. Ghasemi, Y. Fu, X. Ouyang, P. Wang, **M. K. Turkcan**, J. Tavori, S. Kleisarchaki, T. Calmant, L. Gürgen, Z. Kostic, X. Di, G. Zussman, and J. Ghaderi. *ACM/IEEE Symposium on Edge Computing (SEC), 14:1–14:17, 2025.* ***Best Paper Award.*** doi.org/10.1145/3769102.3770618

[C5] **Adaptive Data Collection for Robust Learning Across Multiple Distributions.** C. Zang, **M. K. Turkcan**, G. Zussman, Z. Kostic, and J. Ghaderi. *International Conference on Machine Learning (ICML), PMLR 267:73974–73994, 2025.* proceedings.mlr.press/v267/zang25a.html

[C6] **Loosely Coupled Oscillators as a Correlate of Behavioral Control Circuits Within the Central Complex of the Fruit Fly.** S. Garnell, **M. K. Turkcan**, M. Doborjeh, B. H. Smith, and P. Szyszka. *Neural Information Processing (ICONIP), Lecture Notes in Computer Science, pp. 226–240, 2025.* doi.org/10.1007/978-981-96-6576-1_16

[C7] **StreetNav: Leveraging Street Cameras to Support Precise Outdoor Navigation for Blind Pedestrians.** G. Jain, B. Hindi, Z. Zhang, K. Srinivasula, M. Xie, M. Ghasemi, D. Weiner, S. A. Paris, X. Y. T. Xu, M. Malcolm, **M. K. Turkcan**, J. Ghaderi, Z. Kostic, G. Zussman, and B. A. Smith. *ACM Symposium on User Interface Software and Technology (UIST), pp. 1–21, 2024.* doi.org/10.1145/3654777.3676333

[C8] **Digital Twin for Pedestrian Safety Warning at a Single Urban Traffic Intersection.** Y. Fu, **M. K. Turkcan**, V. Anantha, Z. Kostic, G. Zussman, and X. Di. *IEEE Intelligent Vehicles Symposium (IV), pp. 2640–2645, 2024.* doi.org/10.1109/IV55156.2024.10588544

[C9] **Examining the Influence of Varied Levels of Domain Knowledge Base Inclusion in GPT-based Intelligent Tutors.** B. Castleman and **M. K. Turkcan**. *International Conference on Educational Data Mining (EDM), poster track, pp. 649–657, 2024.* doi.org/10.5281/zenodo.12729908

[C10] **Face-Looking Image Recognition. M. K. Turkcan**, E. Çetin, and T. Akgül. *IEEE Signal Processing and Communications Applications Conference (SIU), pp. 1–4, 2019.* doi.org/10.1109/SIU.2019.8806499

Workshop Papers, Demos, and Short Papers

[W1] **Adapt, Label, Distill: Foundation-Model Supervision for Efficient Object Detection.** Z. Yang, A. Wang, Y. Li, **M. K. Turkcan**, B. Yan, and Z. Kostic. *IEEE/CVF WACV Workshops, Real-World Surveillance Workshop, to appear, 2027.* openreview.net/forum?id=vkrO5Rr8qY

[W2] **Calibration-Free View-Agnostic Monocular 3D Object Detection for Urban Scenes. M. K. Turkcan**, D. Gumaste, and Z. Kostic. *CVPR Workshops (DriveX), pp. 786–795, 2026.* ***Oral presentation.*** openaccess.thecvf.com/CVPR2026_workshops/DriveX

[W3] **Demo: Real-Time Video Analytics for Urban Safety, Deployment over Edge and End Devices.** M. Ghasemi, Y. Fu, X. Ouyang, P. Wang, **M. K. Turkcan**, J. Tavori, S. Kleisarchaki, T. Calmant, L. Gürgen, Z. Kostic, X. Di, G. Zussman, and J. Ghaderi. *ACM/IEEE Symposium on Edge Computing (SEC), demo, pp. 1–3, 2025.* doi.org/10.1145/3769102.3774387

[W4] **Urban Sensing for Human-Centered Systems: A Modular Edge Framework for Real-Time Interaction.** N. Salami Pargoo, M. Ghasemi, S. Xia, **M. K. Turkcan**, T. Ehsan, C. Zang, Y. Sun, J. Ghaderi, G. Zussman, Z. Kostic, and J. Ortiz. *International Workshop on Human-Centered Sensing, Modeling, and Intelligent Systems (HumanSys), at ACM SenSys, pp. 92–97, 2025.* doi.org/10.1145/3722570.3726890

[W5] **Distributed VLMs: Efficient Vision-Language Processing through Cloud-Edge Collaboration.** Y. Li, D. Gumaste, **M. K. Turkcan**, J. Ghaderi, G. Zussman, and Z. Kostic. *IEEE PerCom Workshops, pp. 280–286, 2025.* doi.org/10.1109/PerComWorkshops65533.2025.00078

[W6] **Data-Driven Traffic Simulation for an Intersection in a Metropolis.** C. Zang, **M. K. Turkcan**, G. Zussman, J. Ghaderi, and Z. Kostic. *CVPR Workshops (POETS), 2024.* ***Oral presentation.*** arxiv.org/abs/2408.00943

[W7] **Evaluating Zero-Shot Image Tagging for Urban Traffic Intersection Scenes.** W. Ho, **M. K. Turkcan**, G. Zussman, J. Ghaderi, and Z. Kostic. *CVPR Workshop on Visual Anomaly and Novelty Detection (VAND 2.0), 2024.*

[W8] **Towards Street Camera-based Outdoor Navigation for Blind Pedestrians.** G. Jain, B. Hindi, M. Xie, Z. Zhang, K. Srinivasula, M. Ghasemi, D. Weiner, X. Y. T. Xu, S. A. Paris, C. Tedjo, J. Bassin, M. Malcolm, **M. K. Turkcan**, J. Ghaderi, Z. Kostic, G. Zussman, and B. A. Smith. *ACM SIGACCESS Conference on Computers and Accessibility (ASSETS), pp. 1–6, 2023.* doi.org/10.1145/3597638.3614498

Preprints

[P1] **Dynamic Minimax Regret Optimization for Robust LLM Post-Training.** C. Zang, H. Dong, **M. K. Turkcan**, G. Zussman, Z. Kostic, and J. Ghaderi. *arXiv preprint, 2026.* arxiv.org/abs/2610.06329

[P2] **Reinforcement Learning of Communication in a Mesh of Small Language Models. M. K. Turkcan**. *arXiv preprint, 2026.* arxiv.org/abs/2609.30578

[P3] **Surgical Kinematics from Monocular Video with Learned Articulated Motion Constraints. M. K. Turkcan**, S. Samal, and Z. Kostic. *arXiv preprint, 2026.* arxiv.org/abs/2609.27227

[P4] **Open-H-Embodiment: A Large-Scale Dataset for Enabling Foundation Models in Medical Robotics.** Open-H-Embodiment Consortium: N. Nelson, J.-T. Chen, J. Haworth, et al., including **M. K. Turkcan**. *arXiv preprint, 2026.* arxiv.org/abs/2604.21017

[P5] **A Real-Time Bike-Pedestrian Safety System with Wide-Angle Perception and Evaluation Testbed for Urban Intersections. M. K. Turkcan**. *arXiv preprint, 2026.* arxiv.org/abs/2604.17046

[P6] **Loom: A Scalable Analytical Neural Computer Architecture. M. K. Turkcan**. *arXiv preprint, 2026.* arxiv.org/abs/2604.08816

[P7] **Detect Anything in Real Time: From Single-Prompt Segmentation to Multi-Class Detection. M. K. Turkcan**. *arXiv preprint, 2026.* arxiv.org/abs/2603.11441

[P8] **A Vision-Based Analysis of Congestion Pricing in New York City. M. K. Turkcan**, J. Tavori, J. Ghaderi, G. Zussman, Z. Kostic, and A. Smyth. *arXiv preprint, 2026.* arxiv.org/abs/2602.03015

[P9] **Towards Suturing World Models: Learning Predictive Models for Robotic Surgical Tasks. M. K. Turkcan**, M. Ballo, F. Filicori, and Z. Kostic. *arXiv preprint, 2025.* arxiv.org/abs/2503.12531

[P10] **The Streetscape Application Services Stack (SASS): Towards a Distributed Sensing Architecture for Urban Applications.** N. Salami Pargoo, M. Ghasemi, S. Xia, **M. K. Turkcan**, T. Ehsan, C. Zang, Y. Sun, J. Ghaderi, G. Zussman, Z. Kostic, and J. Ortiz. *arXiv preprint, 2024.* arxiv.org/abs/2411.19714

[P11] **Boundless: Generating Photorealistic Synthetic Data for Object Detection in Urban Streetscapes. M. K. Turkcan**, Y. Li, C. Zang, J. Ghaderi, G. Zussman, and Z. Kostic. *arXiv preprint, 2024.* arxiv.org/abs/2409.03022

[P12] **A Programmable Model for Exploring the Functional Logic of the *Drosophila* Antennal Lobe.** A. A. Lazar, **M. K. Turkcan**, and Y. Zhou. *bioRxiv preprint, 2022.* doi.org/10.1101/2022.09.10.506216

[P13] **A Programmable Model for Exploring the Functional Logic of the *Drosophila* Mushroom Body.** A. A. Lazar, **M. K. Turkcan**, and Y. Zhou. *bioRxiv preprint, 2022.* doi.org/10.1101/2022.09.10.506218

[P14] **The Fruit Fly Brain Observatory: From Structure to Function.** N. H. Ukani, C.-H. Yeh, A. Tomkins, Y. Zhou, D. Florescu, C. Luna Ortiz, Y.-C. Huang, C.-T. Wang, **M. K. Turkcan**, T. Liu, P. Richmond, C.-C. Lo, D. Coca, A.-S. Chiang, and A. A. Lazar. *bioRxiv preprint, 2019.* doi.org/10.1101/580290

[P15] **Using an Ancillary Neural Network to Capture Weekends and Holidays in an Adjoint Neural Network Architecture for Intelligent Building Management.** Z. Ding, **M. K. Turkcan**, and A. Boulanger. *arXiv preprint, 2018.* arxiv.org/abs/1902.06778

Conference Abstracts

[A1] **Optimizing Computer Vision in Surgery Workflows Through Transfer Learning. M. K. Turkcan**, S. Choksi, D. P. Bitner, S. Szot, C. Zang, F. Filicori, and Z. Kostic. *American College of Surgeons Clinical Congress, Boston, MA, 2023. Oral presentation.*

[A2] **Circulating lncRNA ANRIL, a Potential Biomarker in Tanzanian Patients with Coronary Artery Disease.** G. Akan, E. Nyawawa, B. Nyangasa, **M. K. Turkcan**, E. Mbugi, O. Ozgen, M. Janabi, and F. Atalar. *European Journal of Human Genetics, 28(Suppl. 1):264, 2020. ESHG 2020 abstract.* doi.org/10.1038/s41431-020-00739-z

Dissertation

[D1] **Sensory Processing and Associative Learning in Connectome-Based Neural Circuits. M. K. Turkcan**. *Ph.D. dissertation, Columbia University, 2022.* doi.org/10.7916/0h9a-z047

RESEARCH FUNDING

Grants and Computing Awards

◇ NVIDIA research grant, *Generative Edge AI for Dynamic Smart Urban Metropolises* 2026
A. Smyth and **M. K. Turkcan**

◇ NVIDIA research grant, *Edge AI for Equitable and Safe Intersections in Urban Metropolises* 2025
Z. Kostic and **M. K. Turkcan**

◇ EmpireAI award, *AI Computing for Community, Health, Safety and Climate* 2025

Sponsored Research Contributions *Institutional award totals; dates indicate my period of contribution.*

◇ **NSF Engineering Research Center for Smart Streetscapes (CS3)** (EEC-2133516) 2022 – Present
Research staff; research coordination, annual reporting, and sponsor reviews $26M center award
- Coordinate applied AI research and cross-team demonstrations across a five-institution center; present DART, UrbanOmniDetect, and related systems at annual NSF site visits.

◇ **NSF CPS: Hybrid Twins for Urban Transportation** (CNS-2038984) 2021 – Present
Technical, reporting, and sponsor review lead $1.3M Columbia award
- Led engineering delivery, prepared annual technical reports, and presented project reviews to NSF for traffic sensing, prediction, and digital twin systems spanning individual intersections and citywide operations.

◇ **DARPA Robust Learning in Brain Circuits of Synthetic Miniature Insects** 2019 – 2020
Proposal and technical report contributor (HR0011-19-9-0035) $620K Columbia award
- Contributed to proposal development, authored project reports, and conducted grant-supported doctoral research on robust learning and executable brain circuit models.

◇ **AFOSR Foundations of Neuroinformation Processing** (FA9550-16-1-0410) 2017 – 2021
Grant-supported doctoral researcher $1.11M Columbia award
- Developed computational systems and models for phase- and spike-based neural information processing, including work underlying the FlyBrainLab platform.

◇ **NSF Collaborative Research: The Digital Fly Brain** (DBI-1544383) 2017 – 2019
Grant-supported doctoral researcher $795K Columbia award
- Engineered connectome data, graph query, simulation, and visualization capabilities for the open-source Fruit Fly Brain Observatory and FlyBrainLab systems.

◇ **Con Edison Underground Infrastructure and Public Safety Analytics** (CU09-1331) 2015 – 2016
Graduate Researcher; Amendments 21–23 Con Edison
- Developed theoretical models, implementations, and data processing pipelines for a Visual Data Capture pilot using thermal imagery, high-resolution photographs, and video of underground electrical structures; examined automated hotspot detection, inspection targeting, and estimation of avoided serious events.
- Extended the system with infrared camera data assimilation, equipment and cable density models, temperature differential extraction along cable lengths, and visual features associated with later failures; delivered project outputs to Con Edison.
- Contributed cost-benefit and portfolio optimization analysis for secondary public safety programs.

SOFTWARE, MODELS, AND DATASETS

Star and fork counts are from GitHub in November 2026.

Physical AI, Simulation, and Robot Learning

◇ **G1 Patrol** (2026, creator): navigation for a Unitree G1 humanoid through procedurally generated New York City sidewalks, sidewalk sheds, crosswalks, construction zones, doorways, and plazas with simulated pedestrians, plus civic tasks such as litter collection. A privileged PPO teacher is distilled into students that see only through lidar or a depth camera, trained and evaluated in Isaac Lab with domain randomization and tests across the Newton and PhysX simulators. The same pipeline trains a Unitree B2-W quadruped, with a ROS 2 robot stack and a camera encoder aligned to realistic video generated by NVIDIA Cosmos for sim-to-real transfer. Code: github.com/mkturkcan/g1_patrol; project: mehmetkeremturkcan.com/g1_patrol.

◇ **Valdrada** (2026, creator): a real-time digital twin of New York City compiled from public municipal records. It covers all five boroughs with simulated traffic and pedestrians, labels every frame with segmentation, depth, and 3D boxes, can be driven from Python through an API modeled on CARLA's, and exports scenes to OpenUSD for rendering in Blender and Unreal Engine 5. 258 GitHub stars and 37 forks in its first six weeks. Code: github.com/mkturkcan/valdrada; project: mehmetkeremturkcan.com/valdrada.

◇ **Boundless** (2024–2026, creator): an Unreal Engine 5 synthetic data simulator and datasets for infrastructure, aerial, and city-scale object detection with automatic 2D and 3D bounding boxes; detectors trained on its data outperform those trained on CARLA data when tested on real cameras. Code: github.com/mkturkcan/boundless; data: huggingface.co/datasets/mehmetkeremturkcan/boundless-i2x.

◇ **SyntheticStreetScenes** (2026, creator): about 5,200 generated videos of street scenes covering traffic, civic, inspection, and hazard scenarios from infrastructure and egocentric viewpoints, used to train StreetPrompt without class labels. Data: huggingface.co/datasets/mehmetkeremturkcan/SyntheticStreetScenes.

◇ **bikeped** (2026, creator): real-time fisheye perception, collision warnings, simulator scenarios, and conformance testing for urban intersections, with a public dataset. Code: github.com/mkturkcan/bikeped; data: huggingface.co/datasets/mehmetkeremturkcan/bikeped.

Perception and Vision-Language Models

◇ **DART** (2026, creator): real-time open-vocabulary detection from promptable frontier vision models through shared backbone inference, batched decoding, adapter distillation, and TensorRT FP16 optimization. 796 GitHub stars and 93 forks. Code: github.com/mkturkcan/DART; models: doi.org/10.57967/hf/8768.

◇ **UrbanOmniDetect and UrbanOmniDetect-2** (2026, creator): calibration-free monocular 3D detection for ego-vehicle, infrastructure, and aerial cameras; UrbanOmniDetect-2 adds COCO-80 detection in five model scales with real-time tracking and bird's-eye views. Released with the UrbanOmniView multi-view dataset. 136 GitHub stars. Code: github.com/mkturkcan/urbanomnidetect; models: doi.org/10.57967/hf/10583.

◇ **StreetPrompt** (2026, creator): an open-vocabulary video classifier for street cameras that recognizes situations described in plain text, without labeled data or retraining, on hardware from a Jetson Orin Nano to desktop GPUs; trained without class labels on SyntheticStreetScenes and on real traffic footage captioned by a vision-language model. Models: doi.org/10.57967/hf/10582.

◇ **UrbanAnonymizer** (2026, creator): face and license plate anonymization for street-level images and video that outperforms EgoBlur Gen2 on the UFDD and PP4AV benchmarks. Code: github.com/mkturkcan/urbananonymizer; models: doi.org/10.57967/hf/10581.

◇ **Constellation** (2024, dataset lead): a high-altitude urban intersection detection dataset with benchmarks and pretrained models, and a CARLA counterpart rendered from the same camera for sim-to-real evaluation. Code: github.com/zk2172-columbia/constellation-dataset; data: huggingface.co/datasets/mehmetkeremturkcan/constellation_urban_intersection_dataset.

◇ **NYC congestion pricing analysis** (2025–2026, creator): a public dataset and code measuring traffic at 910 public cameras; the Center for Smart Streetscapes runs its real-time camera map. Code: github.com/mkturkcan/congestionpricing; data: huggingface.co/datasets/mehmetkeremturkcan/nyc-congestionpricing-cv.

Medical Robotics

◇ **Surgical SAM 3.1** (2026, with S. Samal and Z. Kostic): a detector and instance segmenter for surgical instruments, anatomy, and tissue that takes text prompts, fine-tuned from SAM 3.1 on four public surgical datasets. Models: huggingface.co/mehmetkeremturkcan/SAM3-Joint-Surgical.

◇ **Suturing world models** (2025, creator): video generation models fine-tuned to predict robotic suturing, with public weights. Code: github.com/mkturkcan/suturingmodels.

Language Models and Agents

◇ **generative-agents** (2023, creator): an open-source implementation of generative agents. 997 GitHub stars and 163 forks. Code: github.com/mkturkcan/generative-agents.

◇ **Compact language and vision-language models** (2025, creator): the SmollerLM family of small language models, trained with supervised fine-tuning, DPO, and pruning, and the FemtoVLM family of small vision-language models. Models: huggingface.co/mehmetkeremturkcan.

◇ **Mote** (2026, creator): runs Gemma-4-E2B at 11 tokens per second on the four CPU cores of a Raspberry Pi 5. Code: github.com/mkturkcan/mote; models: doi.org/10.57967/hf/9204.

◇ **Loom** (2026, creator): an analytical neural computer that runs compiled C programs inside a looped transformer, with ONNX and WebGPU demos, FPGA verification, and 140+ regression tests. Code: github.com/mkturkcan/Loom; models: doi.org/10.57967/hf/8077.

◇ **GPTune** (2019, creator): an early GPT-2 fine-tuning and sampling toolkit for adapting the 774M-parameter model on a single GPU. Code: github.com/mkturkcan/GPTune.

Scientific and Educational Platforms

◇ **FlyBrainLab** (Ph.D. work, co-creator): natural language and ontology interfaces, OrientDB graph querying, GPU simulation, and 3D visualization for connectome-scale neuroscience. 121 GitHub stars. Code: github.com/FlyBrainLab/FlyBrainLab.

◇ **CS3 AI Teaching Kit** (2026, creator): browser lessons and an offline NVIDIA Jetson Orin Nano kit for K–12 classrooms. Kit: huggingface.co/datasets/mehmetkeremturkcan/cs3-ai-teaching-kit.

RESEARCH IMPACT AND ADOPTION

◇ **Citations:** 275 on Google Scholar with an h-index of 12 as of October 2026, rising from 26 citations in 2024 to 77 in 2025 and 132 in 2026 through October. The suturing world models paper is cited by NVIDIA's Cosmos-H-Surgical, and DART by a survey of efficient Segment Anything models in the International Journal of Computer Vision.

◇ **DART:** proposed for integration into Intel's Open Edge Platform with OpenVINO models and into the OCTRON tool for tracking animals; benchmarked against DeepLabCut at the Lemanic Life Sciences Hackathon 2026; extended by DAIRT to detect interactions between objects.

◇ **UrbanOmniDetect:** benchmarked at 30 frames per second on a laptop's integrated GPU by condados.ai and ported to C++ and Qt as OmniBEV.

◇ **Valdrada:** its design, which streams the city in tiles, is followed by Procedural Tokyo and a San Juan port; 3DStreet documented it as prior art, and Databook.nyc asked to link it from its building pages.

◇ **generative-agents:** the base of the Chinese generative agents project in the Luotuo LLM community.

◇ **FlyBrainLab:** its eLife paper has an Altmetric score higher than 87% of papers from the same period.

INVITED TALKS AND PRESENTATIONS

◇ **Valdrada, a Real-Time Digital Twin of New York City** Nov. 10, 2026
Sixth Urban Tech Summit, *The Generative City: New Partnerships for Urban Tech Innovation*, Cornell Tech, New York, NY (scheduled)

◇ **A Vision-Based Analysis of Congestion Pricing in New York City** Sep. 28, 2026
Poster, PAWR PI Meeting 2026, Washington, DC; featured in the PAWR Project Office meeting highlights

◇ **DART and UrbanOmniDetect Research Demonstrations** May 12–15, 2026
CS3 NSF Annual Site Visit, Columbia University

◇ **From Sensors to Systems: Real-Time Traffic Analysis for Faster Decision-Making** Feb. 5, 2026
Panelist, CS3 Innovation Summit, Columbia University, with the NYC Mayor's Office and INRIX

◇ **Port Authority Tomorrow Summit Workshop** Jan. 7, 2026
Workshop host, Port Authority of New York and New Jersey, World Trade Center, New York, NY

◇ **Evaluating Micromobility & Dangerous Riding Behaviors** Nov. 19, 2025
Eighth Annual Vision Zero Research on the Road, New York, NY

◇ **Data-Driven Traffic Simulation for an Intersection in a Metropolis** Jun. 2024
Oral presentation with C. Zang, CVPR 2024 POETS Workshop, Seattle, WA

◇ **Optimizing Computer Vision in Surgery Workflows Through Transfer Learning** Oct. 2023
Oral presentation, American College of Surgeons Clinical Congress 2023, Boston, MA

◇ **FlyBrainLab: An Interactive Open Computing Platform for Exploring the *Drosophila* Brain** Nov. 15, 2020
PyData Global 2020, online

◇ **Drosophila Neurobiology: Genes, Circuits & Behavior Course** Jul. 2018
Lecture on computational approaches for building a functional map of the fruit fly brain, Cold Spring Harbor Laboratory, NY

◇ **Invited Research Talks** 2017 – 2022
Princeton University; University of California, Davis; Max Planck Institute for Biological Intelligence

TEACHING

Columbia University

◇ **Instructor**, ECBM E4040 Neural Networks and Deep Learning, 82 graduate students — Fall 2023

◇ **Instructor**, EECS E6691 Advanced Deep Learning, 41 graduate students — Spring 2023

◇ **Co-instructor**, VOAI 0003E Neural Networks and Deep Learning, Columbia Video Network AI Executive Certificate — Fall 2022, Summer 2023

◇ **Course support**, Computing with Brain Circuits of Model Organisms; Fruit Fly Brain as a Neuroinformation Processor; Computational Neuroscience — 2017 – 2021

◇ EECS E6691 moved from evaluation and segmentation through two-stage and one-stage detection, tracking, transformers, vision-language models, diffusion, and open-vocabulary systems, with a conference-style capstone project. ECBM E4040 covered the mathematical foundations, implementation, and design of neural networks through modern generative models.

◇ Recognized with the Columbia University Electrical Engineering Distinguished Teaching Award (2024).

NSF Center for Smart Streetscapes Research Experience for Teachers

◇ **Main engineering instructor** for three annual cohorts of K–12 STEM educators — 2024 – 2026

◇ Designed and led a four-week engineering curriculum covering neural networks, model evaluation, computer vision, annotation, YOLO training, object tracking, language and generative models, edge deployment, and agents and tool use.

◇ Introduced the CS3 AI Teaching Kit in 2026, pairing browser-based lessons with an offline NVIDIA Jetson Orin Nano kit of 13 camera lessons and 25 sensor lessons; wrote a ten-part Intro to AI lecture series and 12 lecture notes and handouts.

◇ Teachers produced 11 published, classroom-ready lesson plans connecting AI engineering to science, mathematics, civics, and career education (2024 and 2025 cohorts).

Outreach

◇ **Engineering mentor**, CS3 My Streetscape high school research program, on AI models for digital twins of street intersections — 2023

RESEARCH MENTORING

◇ Mentored 11 Columbia Master's students and three high school researchers as a postdoctoral researcher (2022–2024), and undergraduate and high school researchers on FlyBrainLab and the Fruit Fly Brain Observatory during my Ph.D. (2018–2022).

◇ Mentored the student first authors of papers at ICML 2025, IEEE INFOCOM 2026, USENIX VehicleSec 2026, ICMLA 2026, IEEE PerCom Workshops 2025, EDM 2024, Electronics 2025, and Bioengineering 2023.

◇ Six mentees went on to Ph.D. programs at Princeton, the University of Toronto, UC Santa Cruz, and Columbia.

PROFESSIONAL SERVICE

◇ **Poster and Demo Co-Chair**, IEEE SMARTCOMP 2027, 13th IEEE International Conference on Smart Computing, Boca Raton, FL — 2027

◇ **Co-Chair**, CS3 Student Leadership Council, NSF Center for Smart Streetscapes — 2024 – 2025

◇ **Conference reviewer:** CVPR 2026, ICML 2026, ICMLA 2026, WACV 2025, and ACM MobiCom 2024.

◇ **Journal reviewer:** IEEE/ACM Transactions on Networking; IEEE Transactions on Very Large Scale Integration Systems; Sensors, Electronics, and Smart Cities (MDPI).

MEDIA COVERAGE

◇ *Highlights and Materials from the PAWR PI Meeting 2026*, PAWR Project Office — Oct. 2026
advancedwireless.org/outreach/highlights-and-materials-from-the-pawr-pi-meeting-2026/

◇ *How AI Is Improving Daily Life in Cities*, Columbia Engineering — Mar. 2026
www.engineering.columbia.edu/about/news/how-ai-improving-daily-life-cities

◇ *From Smart Streets to Smarter Cities*, INRIX, written with A. Smyth, A. Kosikova, and B. Khalili — Mar. 2026
inrix.com/blog/from-smart-streets-to-smarter-cities-validating-and-scaling-traffic-volume-estimation-in-nyc/

◇ *The Confidence Gap: How to Integrate AI and Traffic Data for Operational Decision-Making*, The Algorithmic Commons — Feb. 2026
algorithmiccommons.substack.com/p/the-confidence-gap-how-to-integrate

◇ Coverage of the SEC 2025 Best Paper Award, Columbia Electrical Engineering — 2025
www.ee.columbia.edu/news/phd-student-mahshid-ghasemi-receives-best-paper-award-acmieee-symposium-edge-computing

◇ *Research Team Showcases AI-based Urban Traffic Applications at CVPR 2024*, Columbia WiMNet Lab — Jul. 2024
wimnet.ee.columbia.edu/research-team-showcases-ai-based-urban-traffic-applications-at-cvpr-2024/

◇ *A Digital Twin Project Using the COSMOS Testbed Named Winner in IDC Smart Cities North America Awards*, COSMOS — 2023
cosmos-lab.org/a-digital-twin-project-using-the-cosmos-testbed-named-winner-in-idc-smart-cities-north-america-awards/

OTHER RESEARCH EXPERIENCE

◇ **Columbia University**, New York, NY — 2015 – 2016
Graduate Researcher
- Developed deep learning systems for thermal camera infrastructure inspection, solar microgrid control, and intelligent building forecasting.

TECHNICAL SKILLS

◇ **AI and deployment:** PyTorch, JAX, TensorFlow, Hugging Face, Ultralytics, OpenCV, TensorRT, ONNX, LiteRT, NCNN, WebGPU; NVIDIA Jetson, Qualcomm, and Raspberry Pi edge devices; detection, segmentation, tracking, 3D perception, VLM/VLA systems, LLM post-training (DPO, RL), RAG, agents, diffusion, and world models.
◇ **Simulation and robotics:** NVIDIA Isaac Lab and Isaac Sim, Newton (MuJoCo-Warp), Unreal Engine 5 (UE5), CARLA, OpenUSD, Blender, ROS 2; humanoid and quadruped robot learning for the Unitree G1 and B2-W (PPO, teacher-student distillation, domain randomization); sim-to-real transfer, city-scale digital twins, traffic and pedestrian simulation, synthetic data, and hardware-in-the-loop testing.
◇ **Languages, systems, and data:** Python, C/C++, CUDA, TypeScript/JavaScript, Node.js, SQL, MATLAB; OrientDB, RPC, distributed querying, HPC, Docker, GitHub Actions; large-scale data collection and annotation (CVAT), OpenStreetMap and NYC Open Data, JupyterLab, Three.js/WebGL, scientific visualization, and Adobe Creative Suite.

CREATIVE WORK

Independent games with Wisedawn (wisedawn.itch.io) and films with KEDIKAT (www.kedikatstudios.com). KEDIKAT's film *FLEA* received four festival accolades in 2026, including Best Storytelling Runner-Up at the MetaMorph AI Award. This work with real-time engines and visual storytelling informs simulation projects such as Boundless.

REFERENCES

Available upon request.